Draft version October 19, 2016
Preprint typeset using LaTeX style emulateapj v. 08/22/09

arXiv:1602.00691v2 [astro-ph.EP] 18 Oct 2016

THE MEASUREMENT, TREATMENT, AND IMPACT OF SPECTRAL COVARIANCE AND BAYESIAN PRIORS IN INTEGRAL-FIELD SPECTROSCOPY OF EXOPLANETS

Johnny P. Greco[1] and Timothy D. Brandt[2,3]

Draft version October 19, 2016

ABSTRACT

The recovery of an exoplanet's atmospheric parameters from its spectrum requires accurate knowledge of the spectral errors and covariances. Unfortunately, the complex image processing used in high-contrast integral-field spectrograph (IFS) observations generally produces spectral covariances that are poorly understood and often ignored. In this work, we show how to measure the spectral errors and covariances and include them self-consistently in parameter retrievals. By combining model exoplanet spectra with a realistic noise model generated from GPI early science data, we show that ignoring spectral covariance in high-contrast IFS data can both bias inferred parameters and lead to unreliable confidence regions on those parameters. This problem is made worse by the common practice of scaling the χ^2 per degree of freedom to unity; the input parameters then fall outside the 95% confidence regions in as many as $\sim$80% of noise realizations. The biases we observe can approach the typical levels of precision achieved in high-contrast spectroscopy. Accounting for realistic priors in fully Bayesian retrievals can also have a significant impact on the inferred parameters. Plausible priors on effective temperature and surface gravity can vary by an order of magnitude across the confidence regions appropriate for objects with weak age constraints; priors for objects with good age constraints are dominated by modeling uncertainties. Our methods are directly applicable to existing high-contrast IFSs including GPI and SPHERE, as well as upcoming instruments like CHARIS and, ultimately, WFIRST-AFTA.

Subject headings: methods: data analysis – techniques: imaging spectroscopy – planetary systems

1. INTRODUCTION

Large ground-based telescopes with adaptive optics have now imaged tens of substellar companions with masses ranging from that of massive brown dwarfs to $\sim$3 $M_{\rm Jup}$ giant planets (e.g., Marois et al. 2008; Thalmann et al. 2009; Kuzuhara et al. 2013; Rameau et al. 2013; Macintosh et al. 2015; Wagner et al. 2016). Direct imaging collects photons emitted by the planets and brown dwarfs, allowing them to be characterized by their colors (Burrows et al. 1997; Knapp et al. 2004; Saumon & Marley 2008; Skemer et al. 2014), and ultimately by their chemistry and composition (Bowler et al. 2010; Barman et al. 2011; Janson et al. 2013; Konopacky et al. 2013). A planet's temperature, mass, and age can be used to infer its formation conditions (Marley et al. 2007; Spiegel & Burrows 2012), while its atmospheric abundances might show where in the protoplanetary disk it accreted its atmosphere (Öberg et al. 2011; Madhusudhan 2012). The field of high-contrast imaging has expanded dramatically in recent years and continues to push to higher contrasts and smaller separations (Bowler 2016, and references therein).

High-contrast imaging has long relied on observing strategies and image processing to achieve a factor of $\sim$10 improvement over an instrument's raw contrast. Surveys with the *Hubble Space Telescope* "rolled" the spacecraft to rotate the point-spread function (PSF) with respect to the field, subtracting the PSF at different orientations (Schneider & Silverstone 2003). Angular differential imaging (ADI) enables the same technique to be used in ground-based imaging (Marois et al. 2006), with additional image processing to model and subtract the stellar PSF (Lafrenière et al. 2007; Soummer et al. 2012). This image processing partially subtracts the flux from faint companions, with the magnitude of the effect varying with position. The partial subtraction is usually estimated by introducing and reducing fake point sources (e.g. Lafrenière et al. 2007), though it may also be forward-modeled as a perturbation to the data reduction (Brandt et al. 2013; Pueyo 2016). An alternative observing strategy known as spectral differential imaging (SDI) simultaneously images a star in two or more neighboring wavelengths. The images are then subtracted to remove starlight; companions with strong spectral features will remain in the residual images (Racine et al. 1999; Marois et al. 2000). SDI has been used in high-contrast surveys with NICI on Gemini South (Biller et al. 2013; Nielsen et al. 2013; Wahhaj et al. 2013) and with NaCo on the VLT (Biller et al. 2007; Maire et al. 2014). Unfortunately, the contrast of SDI imaging is a strong function of the companion's spectrum.

The new generation of high-contrast instruments are integral-field spectrographs (IFSs; e.g., Beuzit et al. 2008; Macintosh et al. 2008; Hinkley et al. 2011; Peters et al. 2012), which promise better performance but pose a new set of data analysis challenges. An IFS simultaneously obtains a spectrum of each spatial element in its field-of-view, producing a three-dimensional data cube composed of one spectral (λ) and two spatial (x, y) dimensions. The stellar PSF can now be modeled and subtracted in three dimensions; image processing can exploit the scaling of diffraction with wavelength to improve the instru-

[1] Department of Astrophysical Sciences, Princeton University, Princeton, NJ, USA
[2] Institute for Advanced Study, Princeton, NJ, USA
[3] NASA Sagan Fellow

ment's raw contrast beyond what is possible with ADI imaging (Sparks & Ford 2002; Berton et al. 2006; Marois et al. 2014). IFSs also naturally enable the recovery of an exoplanet's spectrum. With current image processing techniques, the spectrum is subject to the same partial subtraction biases as high-contrast photometry. It also suffers from spectral covariance. The stellar PSF scales with wavelength but is coherent on spatial scales of λ/D; errors in the model PSF couple neighboring wavelengths in the exoplanet's spectrum.

Exoplanet spectra obtained with IFSs have shown clear signs of clouds and non-equilibrium chemistry in the atmosphere of HR8799 b (Bowler et al. 2010; Barman et al. 2011), as well as the presence of carbon monoxide and water in the atmosphere of HR8799 c (Konopacky et al. 2013). The large grids of substellar model atmospheres now available from several groups (Allard et al. 2011; Spiegel & Burrows 2012) may be used in a global fit to the spectrum (Rice et al. 2015). An alternative approach is to fit the abundances and temperature profiles of the most important species to the observed spectrum (Line et al. 2015). Either approach, using grids of atmospheres that assume equilibrium chemistry and (typically) Solar abundance patterns or by fitting molecular species individually, requires the spectral errors and covariances to be properly modeled and understood.

In the high-contrast regime, it becomes challenging to understand the effects of data analysis techniques on the recovered spectra. The Gemini Planet Imager (GPI) recently discovered a low-mass companion to the young F-type star 51 Eridani (Macintosh et al. 2015). The authors of this study extracted the spectrum of 51 Eri b using three different data analysis pipelines, but were unable to characterize any of their spectra well enough to apply statistical retrieval techniques. The most promising young stars have already been searched for substellar companions with high-contrast photometric surveys (Biller et al. 2013; Chauvin et al. 2015; Brandt et al. 2014a; Bowler et al. 2015). As a result, additional discoveries by GPI and other high-contrast IFSs will likely be at small separations and high contrasts, where the effects of image processing are hardest to characterize. The same challenges will arise in space-based IFS imaging at exceptionally high contrasts using WFIRST-AFTA (Spergel et al. 2015).

In this paper, we measure the spectral covariance in GPI commissioning and early science data. Using our measured spectral covariances, and assuming the data reduction to be unbiased, we demonstrate the effects of incorrect error modeling on atmospheric retrievals. The paper is organized as follows. In Section 2, we define the likelihood function. In Section 3, we discuss sources of spectral covariance in IFS data and show how to measure the spectral covariance matrix using actual IFS observations. In Section 4, we compare results from mock retrievals that use the full covariance matrix, which we assume to be the source of the noise, with retrievals that set its off-diagonal terms to zero. In Section 5, we explore the additional effect of realistic priors on effective temperature and surface gravity in Bayesian parameter retrievals. We conclude in Section 6.

2. THE LIKELIHOOD FUNCTION

Spectroscopic characterization is the process of inferring an object's physical properties from its measured spectrum. A spectroscopic observation generally yields a noisy measurement $\boldsymbol{S}$ of a source's flux density $\boldsymbol{F}$. In the most general case, the measured spectral values may be biased and coupled by a measurement matrix $\boldsymbol{M}$, so that the measured spectrum is given by

$$\boldsymbol{S} = \boldsymbol{M} \cdot \boldsymbol{F} + \boldsymbol{\varepsilon}, \tag{1}$$

where $\boldsymbol{\varepsilon}$ is the matrix of measurement errors. Our goal in this paper is to describe the proper treatment of measurement errors; we assume that $\boldsymbol{M}$ is known (i.e., its inverse can be applied to $\boldsymbol{S}$), and that we can simply write

$$S_i = F_i + \epsilon_i, \tag{2}$$

where the index refers to the i^{th} wavelength bin, and ϵ_i is the measurement error associated with that bin. The value of each ϵ_i is typically impossible to calculate; spectroscopic characterization relies on knowing its statistical properties. A forward model maps the physical quantities of interest (temperature, abundances, etc.) onto theoretical spectra, which may be statistically compared to the observed spectrum $\boldsymbol{S}$ to derive parameter constraints.

Parameter retrieval begins with the likelihood function $\mathcal{L}$, the probability of measuring the observed spectrum $\boldsymbol{S}$ given a model of the true spectrum $\boldsymbol{F}$. The forward model is a function of a set of physical parameters Φ, making $\mathcal{L}$ also a function of Φ. If the probability distribution of ϵ_i in each wavelength bin is Gaussian with zero mean, then the log of the likelihood function is given by

$$-2\ln\mathcal{L}(\Phi) \equiv \chi^2 = (\boldsymbol{S}-\boldsymbol{F})^T \boldsymbol{C}^{-1} (\boldsymbol{S}-\boldsymbol{F}), \tag{3}$$

where $\boldsymbol{C}$ is the covariance matrix with elements

$$C_{ij} \equiv \langle \epsilon_i \epsilon_j \rangle, \tag{4}$$

with $\langle ... \rangle$ denoting the expectation value. In the case of independent errors, $C_{ij} = 0$ for $i \neq j$, and Equation (3) reduces to the usual

$$-2\ln\mathcal{L}(\Phi) = \sum_i \frac{(S_i - F_i)^2}{C_{ii}}. \tag{5}$$

The maximum likelihood estimate is the set of parameters Φ that maximize $\mathcal{L}$ or, equivalently, minimize χ^2. In a Bayesian parameter retrieval, the posterior probability distribution of the parameters Φ given the observed spectrum $\boldsymbol{S}$ is proportional to the product of the likelihood function and the prior probability distribution of Φ:

$$p(\Phi|\boldsymbol{S}) \propto \mathcal{L}(\Phi) p(\Phi). \tag{6}$$

Equation (6) must then be integrated and normalized to obtain a probability density, for example by Markov Chain Monte Carlo. Confidence regions are enclosed by surfaces of constant posterior probability density containing a given fraction of the integrated posterior probability.

3. SPECTRAL COVARIANCE IN IFS DATA

In this paper, we explore the importance of using the correct covariance matrix $\boldsymbol{C}$ to infer physical parameters from IFS data. Many sources of noise, e.g. photon noise and read noise, are independent at different wavelengths,

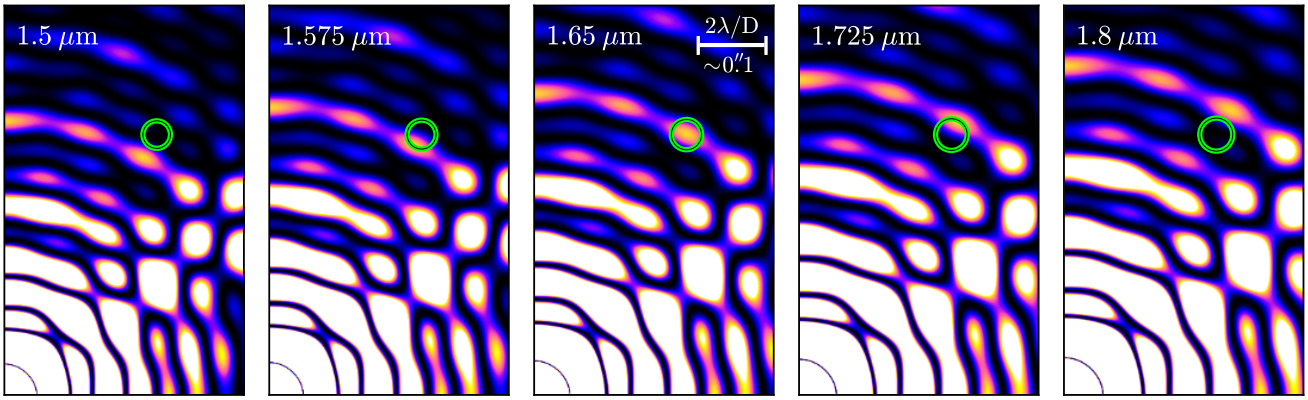

FIG. 1.— A demonstration of the motion of a PSF speckle with respect to the location of a hypothetical planet (green circles). Each frame shows a monochromatic, unocculted PSF with a linear stretch; the center of each PSF defines the lower-left corner of the frame. The bandpass width is $\sim$20%, similar to the H-band. The image scale is given in units of λ/D and arcseconds in the middle frame, assuming $\lambda = 1.65$ μm and $D = 8$ m. The green circles represent a planet located $\sim$11 λ/D from the center of the PSF. The angular diameter of each circle is λ/D, the characteristic scale of a planet PSF. The speckle moves approximately two full planet PSF diameters as we step through the $\sim$20% bandpass, adding flux density within the planet's PSF for roughly half the filter width, which couples neighboring wavelengths at this characteristic spectral resolution.

allowing Equation (5) to be used. The process of reducing high-contrast IFS data, however, introduces sources of noise that are spectrally correlated. As a trivial example, interpolation to align a series of images or to scale a model PSF can couple neighboring wavelengths. As another example, the spectrum of each lenslet on the detector is the convolution of the true spectrum and that lenslet's diffraction pattern, with the additional possibility of diffraction or "crosstalk" from a neighboring lenslet. Spectral correlation can also be induced by subtracting an imperfect model of the stellar PSF. We explore this effect further with a toy model in Section 3.1. We then measure spectral covariance in actual IFS data in Section 3.2.

3.1. *Spectral Covariance from PSF Speckles*

As diffraction phenomena, PSF speckles scale radially with wavelength, moving into and out of a planet's location. This scaling may be exploited to suppress speckles in high-contrast images (Racine et al. 1999; Marois et al. 2000; Sparks & Ford 2002). Imperfectly subtracted speckles will add flux density to a planet's spectrum with a characteristic spectral resolution, which depends upon the angular separation between the planet and star in units of λ/D, where D is the effective telescope diameter. For a change in wavelength $\Delta\lambda$, a speckle's angular distance from the star ρ will vary by $\Delta\rho = \rho\,\Delta\lambda/\lambda$. Thus, for a fixed bandpass width $\Delta\lambda/\lambda$, the rate (with respect to a change in wavelength) at which a speckle moves across the core of a planet's PSF will be proportional to its angular separation from the star. This effect introduces a characteristic spectral correlation length that scales linearly with angular separation.

As an intuitive demonstration of how PSF speckles can induce spectral correlation, we construct a toy model that consists of a series of monochromatic, unocculted PSFs spanning a bandpass of width $\Delta\lambda/\lambda \sim 20\%$ (similar to the H-band). For each wavelength, we generate a PSF by taking the Fourier transform of a uniformly illuminated pupil and squaring the amplitude to get the intensity. We then observe a speckle as it moves across the location of a hypothetical planet. The actual pupil we use, as well as our model's lack of a coronagraph, are not important to us here; the radial scaling of speckles is independent of these details.

Figure 1 shows the expansion of the PSF as we step evenly across the $\sim$20% bandpass in five wavelength frames. The indicated wavelengths span the H-band, and the angular scale given in the middle frame assumes an 8 m primary. The green circles represent a possible location of a planet. Each circle has an angular diameter of λ/D and is located $\sim$11 λ/D from the center of the PSF, which defines the lower-left corner of each frame. At this separation, the radial position of a speckle changes by $\Delta\rho \sim 0.2 \times 11\,\lambda/D \sim 2\lambda/D$ (i.e., 2 planet PSF diameters) across the bandpass. The speckle will add flux density within the planet's PSF for approximately half the filter width, coupling neighboring wavelengths at this characteristic spectral resolution.

3.2. *Measuring the Spectral Covariance Matrix*

Even in a multi-planet image, most of the field-of-view is devoid of planets. If no other astrophysical sources are present (e.g., a disk), spectra of these empty regions are realizations of the spectral noise and can be used to measure both its amplitude and covariance. Here, we demonstrate such measurements using reduced data cubes from the GPI collaboration's First Data Release[4]. We start with the data cubes extracted by the GPI team from the two-dimensional detector readouts and use the GPI data reduction pipeline (Maire et al. 2010; Perrin et al. 2014) to implement the LOCI (Lafrenière et al. 2007, GPI primitive: "ADI with LOCI") and KLIP (Soummer et al. 2012, GPI primitive: "KLIP algorithm Angular Differential Imaging") algorithms for PSF subtraction. Both algorithms operate separately on each wavelength and model the PSF with a number of basis images using least-squares optimization; they differ (in GPI's implementation) in *where* the approximation is computed. LOCI approximates the PSF locally in radially extended wedges, while KLIP performs a separate approximation in a user-supplied number of annuli. With annuli = 0

[4] http://www.gemini.edu/sciops/instruments/gpi/public-data

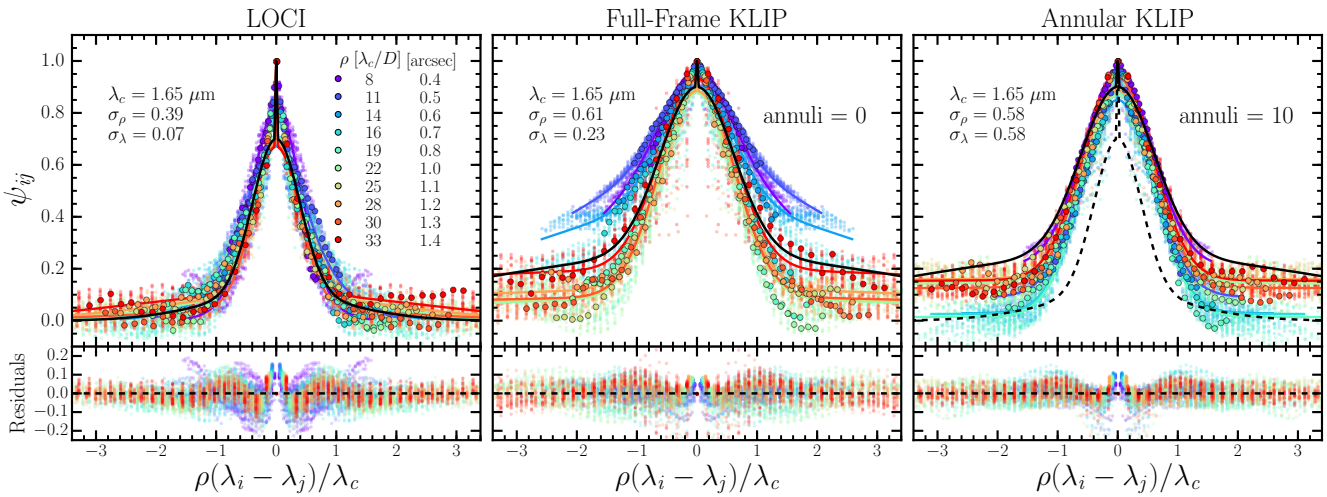

FIG. 2.— The measured spectral correlation ψ_{ij}, as a function of angular separation ρ, in the H-band data cube of HIP 21861, where we have used LOCI (left panel), full-frame KLIP (annuli = 0; middle panel), and KLIP with 10 annuli (right panel) for PSF subtraction. Motivated by the ρ-dependence of PSF-subtraction-induced spectral correlation (see Section 3.1), we plot the correlation functions against $\rho(\lambda_i - \lambda_j)/\lambda_c$, where λ_c is the central wavelength of the bandpass. The colors in each panel correspond to the separations indicated in the left panel in units of λ_c/D and arcseconds (assuming $D = 8$ m). Small dots show the measured correlation for all possible pairs (i, j) of wavelengths, and the larger circles show the correlation for $\lambda_i = \lambda_c = 1.65$ μm (i.e., the middle row of the correlation matrix). The colored lines show the best fits to Equation (8) for each separation, and the lower panels show the associated residuals. Black lines show our fiducial model of the spectral correlation with our LOCI and KLIP noise parameters, which we adopt to perform mock retrievals in Section 4. In the right panel, we show our fiducial models for both KLIP (solid) and LOCI (dashed). LOCI is the most local of these PSF subtraction algorithms and full-frame KLIP is the most global; they bracket a plausible parameter space of spectral covariance.

(the default in the GPI pipeline), KLIP performs a global approximation to the PSF. We adopt LOCI and full-frame (annuli = 0) KLIP as our fidicial reductions. We also compute the spectral covariance for KLIP with annuli = 10; this is more local than full-frame KLIP but less local than LOCI. All other parameters are fixed at GPI's default values (LOCI parameters: nfwhm = 2.5, coeff_type = 0; KLIP parameters: MinRotation = 1°, prop = .99999).

Throughout, we assume the data reduction to be unbiased; for example, the measurement errors are independent of a planet's spectrum. We stress that it is not our goal to fully characterize the noise in GPI data. We instead present a practical approach to dealing with correlated noise in IFS observations and study the general consequences of ignoring spectral covariance when estimating atmospheric parameters from recovered exoplanet spectra. The PSF-subtracted (H-band) data cube of HIP 21861 is by far the cleanest of all the targets in this data release. Since we expect the quality of IFS data will only improve with time, we choose to adopt the HIP 21861 data set as representative of this class of observations. Table 1 summarizes the properties of this observation.

Given a reduced, PSF-subtracted data cube, we measure the average spectral correlation at fixed angular separation ρ from the central star as follows. Assuming azimuthal symmetry, we extract a host-star-centric annulus of radius ρ from each wavelength frame of the data cube. The thickness of the annulus is taken to be 4 pixels (we varied this thickness by ±50% and found that it does not significantly change our results). We then compute the average spectral correlation within the annulus,

$$\psi_{ij} \equiv \frac{C_{ij}}{\sqrt{C_{ii}C_{jj}}} = \frac{\langle I_i I_j \rangle}{\sqrt{\langle I_i^2 \rangle \langle I_j^2 \rangle}}, \quad (7)$$

TABLE 1
THE HIP 21861 GPI DATA SET

Parameter	Value
Observation Date	14 November 2013
Target V-band Brightness	5.0 mag*
Total Time on Target	70 minutes
Total Exposure Time	60 minutes
Field Rotation	95°
Mean Airmass	1.01
Seeing Range	0.″5–0.″7

* From the Tycho-2 catalog (Høg et al. 2000)

where I_i is the intensity at wavelength i, and $\langle ... \rangle$ is the expectation value over all spatial locations within the annulus. Uncorrelated (white) noise would have $\psi_{ij} = 0$ for $i \neq j$.

Figure 2 shows the measured spectral correlation in the H-band data cube of HIP 21861, where we have used LOCI (left panel), full-frame KLIP (annuli = 0; middle panel), and KLIP with 10 annuli (right panel) for PSF subtraction. We plot the correlation functions against $\rho(\lambda_i - \lambda_j)/\lambda_c$, where λ_c is the central wavelength of the bandpass, and ρ is in units of λ_c/D with $D = 8$ m for Gemini. This scaling is motivated by the characteristic correlation length induced by imperfect PSF subtraction, which scales with angular distance from the central star (see Section 3.1). The colors in each panel correspond to the separations indicated in the left panel in units of λ_c/D and arcseconds. The small dots show the measured correlation for all possible pairs (i, j) of wavelengths, and the larger circles show the correlation for $\lambda_i = \lambda_c = 1.65$ μm (i.e., the middle row of the correlation matrix). All panels show spectral correlation on several characteristic scales. In addition, the correlation functions for both LOCI and KLIP show a sharp peak

at the origin, suggesting a significant component of the noise is uncorrelated.

We parameterize the spectral correlation ψ_{ij} with a multi-component noise model consisting of two correlated noise terms and an independent noise term:

$$\psi_{ij} \approx A_\rho \exp\left[-\frac{1}{2}\left(\frac{\rho}{\sigma_\rho}\frac{\lambda_i-\lambda_j}{\lambda_c}\right)^2\right] + A_\lambda \exp\left[-\frac{1}{2}\left(\frac{1}{\sigma_\lambda}\frac{\lambda_i-\lambda_j}{\lambda_c}\right)^2\right] + A_\delta\,\delta_{ij}, \quad (8)$$

where A_ρ and A_λ are the amplitudes of the correlated noise terms, A_δ is the amplitude of the independent noise term, σ_ρ and σ_λ are correlation lengths that characterize the corresponding noise terms, and δ_{ij} is the Kronecker delta. The two correlated noise terms model spectral correlations at different scales. The first term scales with angular separation and models correlations induced from incomplete PSF subtraction. The second term is independent of angular separation and models other sources of spectral correlation, such as interpolation or spectral crosstalk, which do not scale with angular distance from the central star. The Kronecker delta only contributes to the diagonal of the correlation matrix, allowing a component of the noise to be uncorrelated. When fitting for these parameters, we allow the three amplitudes to vary with angular separation but force σ_ρ and σ_λ to be constant across the field-of-view. In an actual parameter retrieval, it may be more appropriate to restrict the fits to a few separations near the planet's location rather than the entire field-of-view, but our aim here is to gain some intuition for the problem and identify potential sources of spectral covariance. In addition, we require the amplitudes to sum to unity. The colored lines in Figure 2 show the best fits for each separation, and the lower panels show the associated residuals.

We note that, given the two correlation lengths (σ_ρ and σ_λ), fitting for the amplitudes is a linear problem. Therefore, it is computationally inexpensive to find the best-fit amplitudes; in this case, sampling algorithms such as Markov Chain Monte Carlo are not necessary. By calculating the best-fit parameters over a sufficiently large (two-dimensional) grid of correlation lengths, one can be certain that they have found the global minimum of the fitting function.

For LOCI, we find typical amplitudes of $A_\rho \sim 0.6$, $A_\lambda \sim 0.1$, and $A_\delta \sim 0.3$, with best-fit correlation lengths $\sigma_\rho = 0.39$ and $\sigma_\lambda = 0.07$. For full-frame KLIP, we find typical amplitudes of $A_\rho \sim 0.65$, $A_\lambda \sim 0.25$, and $A_\delta \sim 0.1$, with best-fit correlation lengths $\sigma_\rho = 0.61$ and $\sigma_\lambda = 0.23$. We refer to these sets of parameters as our LOCI and KLIP noise parameters, respectively, and adopt them as our fiducial model of the spectral correlation in the following sections. The black lines in Figure 2 show these parameterizations of the spectral correlation, which combine with Equation (7) to provide a realistic model of the covariance matrix $\boldsymbol{C}$. For comparison, KLIP with 10 annuli produces typical amplitudes of $A_\rho \sim 0.8$, $A_\lambda \sim 0.1$, and $A_\delta \sim 0.1$, with best-fit correlation lengths $\sigma_\rho = 0.58$ and $\sigma_\lambda = 0.58$. In the right panel of Figure 2, the black lines show the correlation function assuming our fiducial KLIP (solid) and LOCI (dashed) parameters defined above. Generally speaking, the spectral covariance for KLIP with annuli = 10 is intermediate between that for LOCI and for full-frame KLIP.

The spectral correlations for both LOCI and KLIP have a dominant ($A_\rho \gtrsim 60\%$) noise component that is correlated on a scale of $\sigma_\rho \sim 0.5$ in units of λ_c/D. This is consistent with our expectation from Section 3.1 of spectral correlations induced by imperfect PSF subtraction. Both LOCI and KLIP also have a significant component of uncorrelated noise, with LOCI having more power in this term. The most significant difference between LOCI and full-frame KLIP is associated with the second noise term in Equation (8); KLIP produces much stronger large-scale correlations, whereas LOCI appears to be more sensitive to small-scale coupling between neighboring wavelengths.

LOCI and KLIP are mathematically similar and, in GPI's implementation, both operate separately on each wavelength. The reasons for the striking differences in their covariance matrices are subtle. LOCI calculates its PSF approximation locally within radially extended wedges; this radial extent means that LOCI has some knowledge of speckles that, as shown in Figure 1, enter the subtraction region at shorter wavelengths. In contrast, GPI's implementation of KLIP is either a global (annuli = 0) or strictly annular algorithm. At a given wavelength, there is less effect from speckles that enter the subtraction region at other wavelengths. Global fits may also have larger residuals due to the need to fit regions of high intensity, which are disproportionately weighted in least squares. These residuals may have power on many spatial scales. Implementations of KLIP using different geometries, for example radially extended wedges or with smaller annuli, might produce very different spectral covariances. This highlights that the spectral covariance is highly sensitive to the details of how the PSF subtraction algorithm is implemented. It is essential to measure the spectral covariance for each data set and reduction technique. In the GPI pipeline, LOCI and full-frame KLIP (annuli = 0) generate spectral covariance using the most local and global PSF approximations, respectively; we use these for our examples in the rest of this paper.

4. IMPACT ON INFERRED PARAMETERS

We now combine our realistic noise model with theoretical spectra to investigate the effect of spectral covariance on atmospheric parameter retrievals. Assuming the solar-metallicity BT-Settl spectral models (Allard et al. 2011) as "truth", we select a particular spectrum, add a realization of noise using our measured spectral errors and covariances, and attempt to recover the model's input effective temperature ($T_{\rm eff}$) and surface gravity ($\log g$). Only by guaranteeing that the true spectrum is in our fitting library may we satisfy the assumptions of Section 2. We perform thousands of parameter retrievals with and without accounting for the full covariance matrix that we assume to be the source of the noise (i.e., using the full covariance matrix and setting the off-diagonal terms to zero). We run our analysis using both our LOCI and KLIP noise parameterizations. In Section 4.2, we carry out the retrievals with uniform priors in $T_{\rm eff}$ and $\log g$. In Section 5, we study the additional

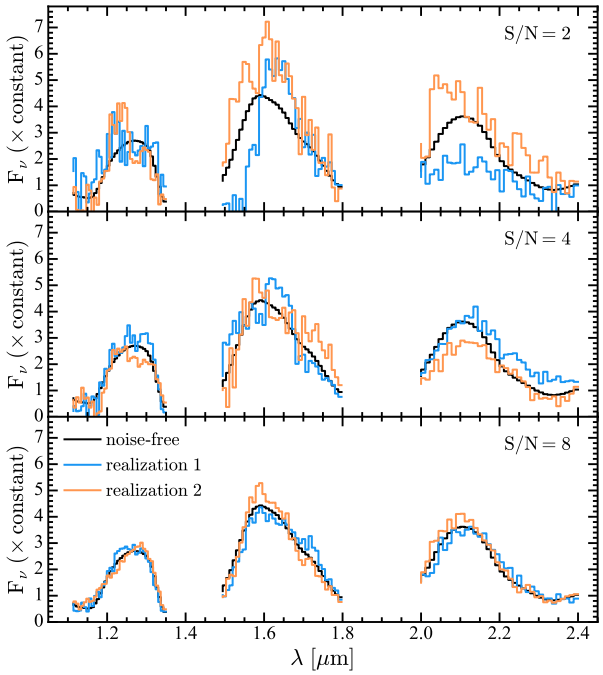

FIG. 3.— Mock spectra of a planet with $T_{\rm eff} = 1000$, $\log g = 4.0$, and solar metallicity generated from the BT-Settl spectral models (Allard et al. 2011). All spectra are smoothed to a constant resolution of $\lambda/\Delta\lambda = 45$ using a Gaussian line spread function and resampled with 37 wavelength bins in J, H, and K. The black lines show the noise-free spectrum of the planet, and the blue and orange lines show different realizations of the spectrum with correlated Gaussian noise (generated with our LOCI parameterization), assuming a mean S/N of 2 (top), 4 (middle), and 8 (bottom). For each band, we define the mean S/N to be the mean flux across the band divided by the root-mean-square of the noise.

effect of realistic priors in Bayesian parameter retrievals.

4.1. *Mock Spectra Generation*

We generate mock spectra from the BT-Settl spectral models; these are provided on grids of effective temperature, surface gravity, and metallicity. We interpolate between models using piecewise power laws and smooth each spectrum to a constant resolution of $\lambda/\Delta\lambda = 45$ using a Gaussian line spread function with full width at half maximum FWHM $= \Delta\lambda/\lambda$. We then resample in J, H, and K with 37 wavelength bins per band, which is similar to the wavelength sampling in the GPI data cubes and somewhat higher than the Nyquist sampling rate. Finally, to each mock spectrum, we add Gaussian noise with spectral covariance given by our LOCI and KLIP noise parameterizations. For the variances C_{ii}, we use the measured variances at a separation of $0\farcs7 \sim 16\,\lambda_c/D$ in the H-band data cube of HIP 21861. Although we only have measurements in $H(1.65\ \mu\text{m}, 16\,\lambda_c/D)$, we use the same (wavelength-dependent) noise parameterizations for $J(1.23\ \mu\text{m}, 22\,\lambda_c/D)$ and $K(2.1\ \mu\text{m}, 12\,\lambda_c/D)$. We note that it is likely that measurements in each band would produce somewhat different noise parameterizations. In an actual parameter retrieval, the covariance matrix should be measured in each band.

Figure 3 shows the noise-free spectrum (black lines) of a planet with $T_{\rm eff} = 1000\ K$, $\log g = 4.0$, and solar metallicity along with two realizations of the spectrum with correlated Gaussian noise (generated with our LOCI noise parameterization), assuming a mean signal-to-noise (S/N) of 2 (top panel), 4 (middle panel), and 8 (bottom panel) in each band. Here, we define S/N to be the mean flux across each band divided by the root-mean-square of the noise. Correlated noise can significantly alter the shape of broad features (e.g., the triangular shape in the H-band), which can impact parameters inferred from such data.

4.2. *Inferring $T_{\rm eff}$ and $\log g$*

To explore how spectral covariance influences parameter estimation, we perform thousands of retrievals of $T_{\rm eff}$ and $\log g$ using the standard method of χ^2-minimization on a grid of forward models, which is widely used in the literature (e.g., Saumon et al. 2006; Cushing et al. 2008; Rice et al. 2010; Madhusudhan et al. 2011; Bonnefoy et al. 2014; Rice et al. 2015). Our approach may be naturally extended to the inverse methods for atmospheric retrievals that are actively being developed for exoplanetary studies (e.g., Line et al. 2013; Lee et al. 2013; Line et al. 2015). In all our calculations, we assume solar metallicity.

For each retrieval, we generate a mock spectrum as described in Section 4.1, adding a realization of spectral noise with S/N=10 to a BT-Settl model. We then build grids of $\chi^2 = -2\ln\mathcal{L}$ using both Equations (3) and (5). We define the former to be the "C case", since the retrieval is performed with the full covariance matrix, and the latter to be the "diag(C) case", since the off-diagonal elements of the covariance matrix are set to zero. The best-fit parameters are then given by the maximum likelihood or, equivalently, the minimum χ^2. For each ($T_{\rm eff}$, $\log g$) pair, we allow the normalization of the spectrum to vary as an additional free parameter, resulting in $37 \times 3 - 3 = 108$ degrees of freedom. Allowing the normalization to float accounts for the generic difficulty models have with producing the correct luminosity in J, H, and K.

Figure 4 shows example grids of $\Delta\chi^2 \equiv \chi^2 - \chi^2_{\rm min} = -2\ln(\mathcal{L}/\mathcal{L}_{\rm max})$ in the $T_{\rm eff} - \log g$ plane for a high-temperature ($T_{\rm eff} = 1810$ K; left column) and low-temperature ($T_{\rm eff} = 750$ K; middle column) input spectrum, where we have modeled the noise with our fit to the LOCI spectral covariance. In the right column, we show the low-temperature case using our KLIP noise parameters. In the top row, we include spectral covariance in our calculations of $\Delta\chi^2$; we ignore it in the middle and bottom rows. To generate each column, we use a single realization of noise; differences between the top and middle rows are due solely to whether or not spectral covariance was included in the fitting procedure. The input model's parameter values are indicated in each panel by yellow stars, and the best-fit parameters are shown as red circles. The solid, dashed, and dotted lines show the 68%, 95%, and 99% confidence regions, respectively. These regions are determined by integrating the likelihood function, which is equivalent to a Bayesian analysis with flat priors in $T_{\rm eff}$ and $\log g$. In the bottom row, we show the same $\Delta\chi^2$ grids as the middle row, except rescaled so that the χ^2 per degree of freedom ($\chi^2_{\rm dof}$) is unity. In all but one case shown in Figure 4, ignoring spectral covariance leads to the true parameter value falling outside the 95% confidence region. Furthermore, rescaling so that $\chi^2_{\rm dof} = 1$ shrinks the confidence regions,

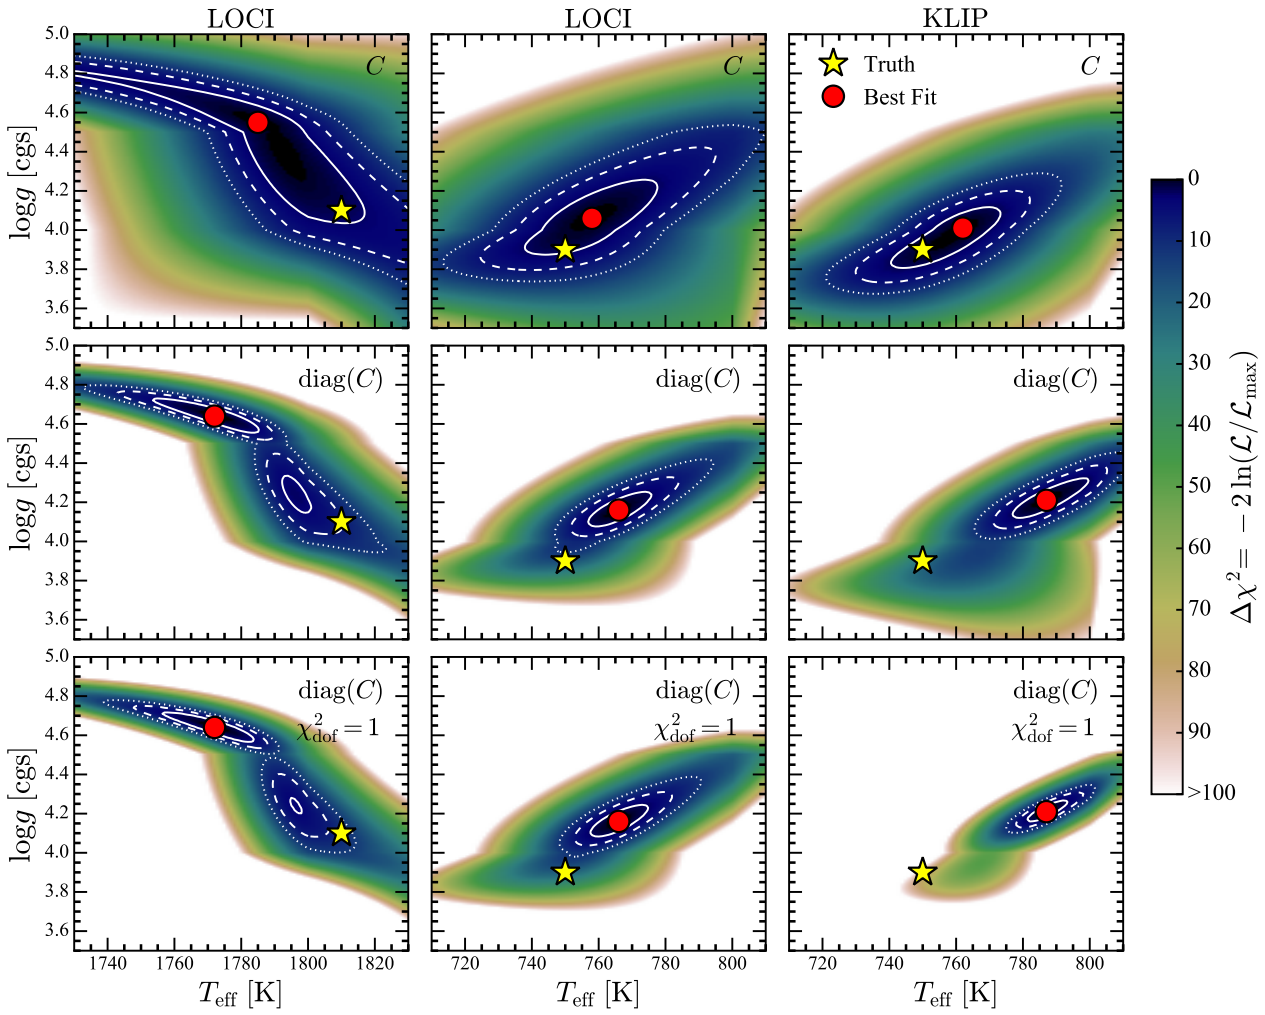

FIG. 4.— Representative grids of $\Delta\chi^2 \equiv \chi^2 - \chi^2_{\rm min} = -2\ln(\mathcal{L}/\mathcal{L}_{\rm max})$ for a high-temperature ($T_{\rm eff}$ = 1810 K; left column) and low-temperature ($T_{\rm eff}$ = 750 K; middle column) input spectrum, assuming our LOCI noise parameterization. The right column shows the low-temperature case assuming our KLIP noise parameterization. In the top row, we calculate χ^2 using the full covariance matrix (C case; Equation (3)), and in the middle and bottom rows, we use the diagonal of the covariance matrix (diag(C) case; Equation (5)). Yellow stars indicate the model's input parameter values, and red circles indicate the best-fit parameters. The solid, dashed, and dotted lines show the 68%, 95%, and 99% confidence regions, respectively, which are determined by integrating $\mathcal{L}$. In the bottom row, we show the same $\Delta\chi^2$ grids as in the middle row, except rescaled so that the χ^2 per degree of freedom ($\chi^2_{\rm dof}$) is unity. We used the same realization of noise to generate each column; differences between the top and middle rows are due solely to whether or not spectral covariance was included in the fitting procedure.

which exacerbates the problem.

Although this figure only shows results for a few realizations of noise, it is representative of the differences we generally see between retrievals that ignore spectral covariance and those that account for it. For both our LOCI and KLIP noise parameters, retrievals that ignore covariance tend to produce confidence regions that are tighter than those that include it. More importantly, the confidence regions in retrievals that ignore covariance are unreliable. In 5×10^4 realizations, the "true" parameter values in such retrievals fell within the 95% confidence regions only ∼33% of the time for LOCI and only ∼23% of the time for KLIP. When $\chi^2_{\rm dof}$ is set to unity, these fractions become ∼ 26% and ∼14%, respectively. The biases in the inferred parameters are generally smaller than the claimed uncertainties for many substellar companions, but are larger than the uncertainties in our simulated retrievals. As uncertainties and systematics in spectral modeling fall, systematics in retrieval techniques will become increasingly important.

In Figure 5, we show the χ^2 distributions from 5×10^4 retrievals of $T_{\rm eff}$ and $\log g$, where we use the full covariance matrix (red histograms) and set its off-diagonal terms to zero (blue histograms). The solid (dashed) lines show the distributions assuming our LOCI (KLIP) noise parameters. The solid black line shows the theoretical expectation for χ^2 with Gaussian noise and 108 degrees of freedom (χ^2_{108}). As expected, we recover the appropriate χ^2 distribution with the correct error model (C case). Ignoring the off-diagonal elements of the covariance matrix (diag(C) case), however, tends to yield χ^2 values that are systematically lower, with the offset being more pronounced for our KLIP noise parameters. This result has important implications, as it means that simply rescaling the uncertainties so that $\chi^2_{\rm dof} = 1$ will generally shrink the confidence regions, making the true parameter values

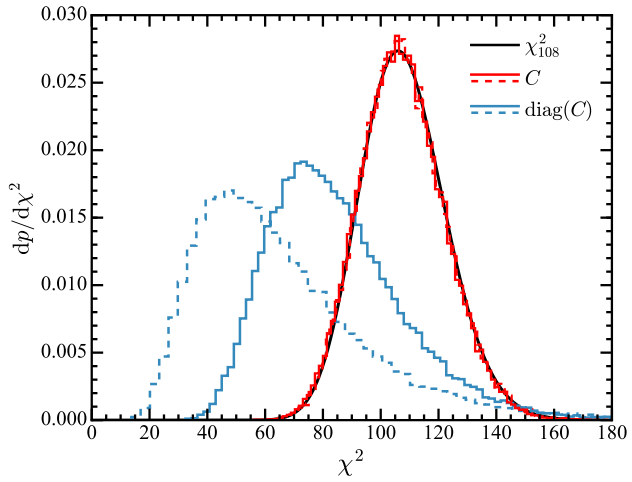

FIG. 5.— χ^2 distributions from 5×10^4 retrievals of $T_{\rm eff}$ and $\log g$, where we calculate χ^2 using the full covariance matrix (red histograms; Equation (3)) and the diagonal of the covariance matrix (blue histograms; Equation (5)). The solid (dashed) lines show the distributions assuming our LOCI (KLIP) noise parameters. The solid black line shows the theoretical expectation for χ^2 with Gaussian noise and 108 degrees of freedom (χ^2_{108}). As expected, we recover the appropriate χ^2 distribution with the correct error model. However, ignoring spectral covariance tends to produce $\chi^2 < 108$, the expectation value of χ^2_{108}, with the offset being more pronounced for KLIP.

fall outside the 95% contours 75% − 85% of the time (c.f. bottom row of Figure 4).

5. THE EFFECT OF REALISTIC PRIORS

In the parameter retrievals shown in Figure 4, we simply mapped the full shape of the likelihood function. This approach is equivalent to a Bayesian analysis when $T_{\rm eff}$ and $\log g$ have uniform priors. However, given assumptions about the mass and age distributions of substellar companions and a substellar cooling model, the prior probability distribution in the $T_{\rm eff} - \log g$ plane ($dp/d\log g/dT_{\rm eff}$) will not, in general, be uniform. Here, we make simple assumptions to calculate a plausible prior $dp/d\log g/dT_{\rm eff}$, showing the effect nonuniform priors can have on recovered $T_{\rm eff}$ and $\log g$ values.

We start by assuming flat priors on the age and logarithmic mass distributions of substellar companions: $dp/dt/d\log M \sim$ constant. Our assumption about the mass distribution is roughly consistent with current observations (e.g., Cumming et al. 2008; Brandt et al. 2014b). The flat prior on age is representative of the common observational scenario in which upper and lower limits on the system's age are all that is known. In this simplified picture, the prior is given by the Jacobian:

$$\frac{dp}{dT_{\rm eff}\, d\log g} \propto \left| \frac{d\log M}{dT_{\rm eff}} \frac{dt}{d\log g} - \frac{d\log M}{d\log g} \frac{dt}{dT_{\rm eff}} \right|. \quad (9)$$

To calculate $dp/dT_{\rm eff}/d\log g$, one can either use Equation (9) directly or Monte Carlo sample $dp/dt/\log M$ and change coordinates into the $T_{\rm eff} - \log g$ plane. We choose the latter, as it avoids numerical issues associated with singularities in the calculation of the Jacobian. We use the solar-metallicity BT-Settl models (Baraffe et al. 2003; Allard et al. 2011) to relate mass and age to $T_{\rm eff}$ and $\log g$. We interpolate between models using second-order splines in the logarithm of all the parameters, and we occasionally linearly extrapolate (in the logarithm of the parameters) the grid at low temperatures ($\lesssim 300$ K)

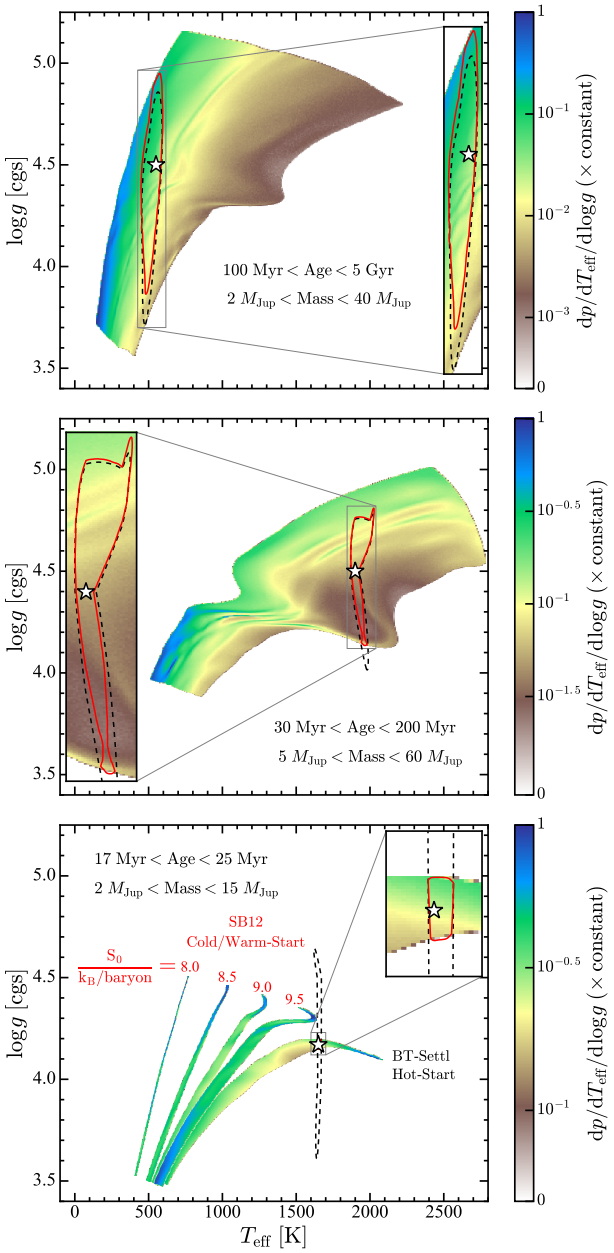

FIG. 6.— Prior probability distributions, $dp/dT_{\rm eff}/d\log g$, assuming flat priors on the age and logarithmic mass distributions of substellar companions. For the substellar cooling model, we use the solar-metallicity BT-Settl models (Baraffe et al. 2003; Allard et al. 2011). The indicated age and mass ranges were chosen to be representative of current estimates for GJ 504b (top panel), κ And b (middle panel), and β Pic b (bottom panel). We perform mock parameter retrievals on the $T_{\rm eff}$ and $\log g$ values indicated by the white stars in each panel. The dashed black and solid red contours show the 95% confidence regions associated with each retrieval, where we have ignored and included the effect of the prior, respectively. The insets magnify where the contours and priors overlap. Each distribution has been normalized such that its peak is unity. In the bottom panel, we also show priors for a series of Spiegel & Burrows (2012, SB12) models with the indicated initial entropies (S_0) in units of Boltzmann's constant (k_B) per baryon. When the age is well-constrained, the prior is dominated by evolutionary model uncertainties.

for masses $\lesssim 6\ M_{\rm Jup}$, with most of the extrapolation being necessary for masses in the range $2-4\ M_{\rm Jup}$.

When characterizing substellar companions with direct-imaging observations, the system's age is one of the most important parameters. The reason for this is that, at fixed luminosity, age and mass are highly degenerate. In the limit of perfect knowledge of a companion's age, the prior $dp/dT_{\rm eff}/d\log g$ will be a line in the $T_{\rm eff}-\log g$ plane (an isochrone). In this case, uncertainties in evolutionary models will dominate the priors. As the age becomes increasingly uncertain, the prior is spread over an increasing range of temperatures and gravities, with the details of the distribution being determined by the physics of substellar cooling and the underlying mass distribution.

The age of the substellar companion κ And b is uncertain by nearly an order of magnitude, falling in the range $\sim 30-200$ Myr (Carson et al. 2013; Hinkley et al. 2013; Bonnefoy et al. 2014). Even more uncertain, age estimates for GJ 504b range from as young as $\sim$100 Myr (Kuzuhara et al. 2013) to as old as $\sim$5 Gyr (Fuhrmann & Chini 2015). These age constraints lead to broad, but far from uniform, priors on $T_{\rm eff}$ and $\log g$. In Figure 6, we show $dp/dT_{\rm eff}/d\log g$ for observations of a system with age constraints similar to κ And b ($30-200$ Myr, middle panel) and GJ 504b (100 Myr $-$ 5 Gyr, top panel). The mass ranges we assume for each calculation, which are indicated in the figure, also span likely ranges for these companions. The peak probability of each distribution is normalized to unity. In both panels, dark blue and white correspond to the peak and zero probability, respectively, but the color scales are otherwise different.

The ripple-like features that are present in the top and middle panels of Figure 6 are due to the onset of deuterium burning. Generically, low-temperature objects have the highest prior probability. We also see that the prior probability of $\log g$ at fixed $T_{\rm eff}$ generally increases with the value of $\log g$. Thus, if the effective temperature of a substellar companion is all that is known, it is more likely to have a high surface gravity. This trend is the result of the nonuniform cooling rates of substellar objects as a function of age (objects cool more rapidly at early times). As $\log g$ (at fixed temperature) increases, objects tend to fall on older isochrones. Since we assume $dp/d\log M\sim$ constant and the radii of these objects are only weak functions of mass and age, we have $dp/d\log g\sim$ constant. Therefore, the prior probability of $\log g$ at fixed temperature depends primarily on how fast objects cool past that temperature, resulting in prior probabilities that increase with $\log g$.

The priors shown in the top and middle panels of Figure 6 both assume poor age constraints. In the bottom panel, we contrast this assumption with an example of a tight age constraint. This age range ($17-25$ Myr) is consistent with current estimates (Binks & Jeffries 2014; Mamajek & Bell 2014) for the substellar companion β Pic b (Lagrange et al. 2009, 2010) and is more representative of the current sample of imaged companions including the HR8799 planets (Marois et al. 2008, 2010), HD 95086 b (Rameau et al. 2013), 51 Eri b (Macintosh et al. 2015), and HD 131399Ab (Wagner et al. 2016). With such a well-constrained age, the details of the model become very important. Therefore, in addition to the BT-Settl "hot-start" model, we show priors for a series of Spiegel & Burrows (2012, SB12) "cold/warm-start" models. We assume SB12's solar-metallicity "hybrid cloud" models; we also performed the calculations at $3\times$ solar metallicity and with their cloud-free models and found the priors to be dominated by the initial entropy. At such a young and precise age, uncertainties in the initial conditions dominate the priors.

To demonstrate the effect of nonuniform priors on recovered $T_{\rm eff}$ and $\log g$ values, we perform retrievals on the "true" parameter values indicated as white stars in Figure 6. These parameters are consistent with current estimates for GJ 504b, κ And b, and β Pic b. We include spectral covariance in each retrieval, and for the noise model, we use our LOCI noise parameters. Additionally, we assume S/N = 10 in the bottom and middle panels and S/N = 4 in the top panel—again, representative of plausible values for observations of these companions. The dashed black contours show the 95% confidence region associated with each retrieval, where we have ignored the effect of the prior (i.e., using only the likelihood function). The solid red contours show the 95% confidence region in a Bayesian parameter retrieval, which includes the effect of the prior (i.e., using the posterior probability distribution). The insets magnify where the contours and priors overlap, clarifying the probability variation within these regions.

In the top panel of Figure 6, where we have assumed an age constraint similar to estimates for GJ 504b, the prior probability varies by about an order of magnitude within the 95% confidence contours. In this case, including the prior shifts the confidence region up in $\log g$ by $\sim$0.1 dex. In the middle panel, where we have assumed age constraints consistent with κ And b, the prior probability varies by a factor of $\sim$5 within the 95% confidence contours. The dominant effect of including the prior in this case is to sharply cut off the confidence region at the edge of the prior's boundary. Uncertainties in evolutionary models dominate the priors shown in the bottom panel, where we have assumed a well-constrained age similar to β Pic b. In this case, our mock retrieval confirms the finding of previous studies that β Pic b's combination of age, surface gravity, and effective temperature is incompatible with cold-start initial conditions (Lagrange et al. 2010; Quanz et al. 2010).

6. SUMMARY AND CONCLUSIONS

In this paper, we have shown how to measure spectral errors and covariances in IFS data and demonstrated the importance of using the full covariance matrix, as opposed to assuming independent errors, in atmospheric parameter retrievals. By measuring the spectral errors and covariances in GPI early science data, we generate a realistic noise model with parameterizations for data PSF-subtracted with GPI's implementation of the LOCI and KLIP algorithms. We find that KLIP—whether its PSF approximation uses the full frame or 10 annuli—produces stronger large-scale spectral correlation than LOCI, which is likely due to subtle differences in the implementations of these algorithms (see Section 3.2). This finding highlights the importance of measuring the spectral covariance for each data set and reduction technique.

We combine our noise model with theoretical exoplanet spectra (Allard et al. 2011) to perform mock retrievals of $T_{\rm eff}$ and $\log g$ with and without the full covariance

matrix that we assume to be the source of the noise. In 5×10^4 realizations, the "true" parameter values in retrievals that ignore spectral covariance fell within the 95% confidence region only ∼33% of the time for LOCI and ∼23% and of the time for KLIP. Scaling $\chi^2_{\rm dof}$ to unity worsens the problem, decreasing these fractions to ∼26% and ∼14%, respectively.

Finally, we explore the additional effect of nonuniform priors on recovered $T_{\rm eff}$ and $\log g$ values. Assuming flat priors on the age and logarithmic mass distributions of substellar companions, we generate prior probability distributions in the $T_{\rm eff} - \log g$ plane. Low-temperature companions have the highest prior probability, and the prior probability of a particular $\log g$ at fixed $T_{\rm eff}$ generally increases with the value of $\log g$. Thus, if the effective temperature of a substellar companion is all that is known, it is more likely to have a high surface gravity. We perform mock Bayesian parameter retrievals including spectral covariance on systems with and without well-constrained ages. For systems with well-constrained ages, we find that the prior probability can vary within the 95% confidence region by as much as an order of magnitude. When the age is tightly constrained, the prior is dominated by uncertainties in evolutionary models.

The new generation of high-contrast instruments are all IFSs, from P1640 (Hinkley et al. 2011), GPI (Ingraham et al. 2014), SPHERE (Beuzit et al. 2008), and CHARIS (Peters-Limbach et al. 2013; Groff et al. 2014) on large ground-based telescopes, to WFIRST-AFTA (Spergel et al. 2015) in space. These new instruments will discover and obtain low-resolution spectra of faint exoplanets close to their host stars, where they would be undetectable without image processing. While correlated noise is inevitable in such spectra, the methods we have presented enable it to be statistically characterized and properly included in Bayesian retrievals of atmospheric parameters.

The authors would like to thank the anonymous referee for a very helpful review that led to significant improvements in this paper. This research is based on observations obtained at the Gemini Observatory, which is operated by the Association of Universities for Research in Astronomy, Inc., under a cooperative agreement with the NSF on behalf of the Gemini partnership: the National Science Foundation (United States), the National Research Council (Canada), CONICYT (Chile), the Australian Research Council (Australia), Ministério da Ciência, Tecnologia e Inovação (Brazil) and Ministerio de Ciencia, Tecnología e Innovación Productiva (Argentina). J.P.G. is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1148900. This work was performed in part under contract with the Jet Propulsion Laboratory (JPL) funded by NASA through the Sagan Fellowship Program executed by the NASA Exoplanet Science Institute. J.P.G. would like to thank Tim Morton for useful conversations. The authors thank Ed Turner and Laurent Pueyo for helpful comments on the manuscript.

REFERENCES

Allard, F., Homeier, D., & Freytag, B. 2011, in Astronomical Society of the Pacific Conference Series, Vol. 448, 16th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun, ed. C. Johns-Krull, M. K. Browning, & A. A. West, 91
Baraffe, I., Chabrier, G., Barman, T. S., Allard, F., & Hauschildt, P. H. 2003, A&A, 402, 701
Barman, T. S., Macintosh, B., Konopacky, Q. M., & Marois, C. 2011, ApJ, 733, 65
Berton, A., Gratton, R. G., Feldt, M., et al. 2006, PASP, 118, 1144
Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2008, in Proc. SPIE, Vol. 7014
Biller, B. A., Close, L. M., Masciadri, E., et al. 2007, ApJS, 173, 143
Biller, B. A., Liu, M. C., Wahhaj, Z., et al. 2013, ApJ, 777, 160
Binks, A. S., & Jeffries, R. D. 2014, MNRAS, 438, L11
Bonnefoy, M., Currie, T., Marleau, G.-D., et al. 2014, A&A, 562, A111
Bowler, B. P. 2016, ArXiv e-prints, 1605.02731
Bowler, B. P., Liu, M. C., Dupuy, T. J., & Cushing, M. C. 2010, ApJ, 723, 850
Bowler, B. P., Liu, M. C., Shkolnik, E. L., & Tamura, M. 2015, ApJS, 216, 7
Brandt, T. D., Kuzuhara, M., McElwain, M. W., et al. 2014a, ApJ, 786, 1
Brandt, T. D., McElwain, M. W., Turner, E. L., et al. 2013, ApJ, 764, 183
Brandt, T. D., McElwain, M. W., Turner, E. L., et al. 2014b, ApJ, 794, 159
Burrows, A., Marley, M., Hubbard, W. B., et al. 1997, ApJ, 491, 856
Carson, J., Thalmann, C., Janson, M., et al. 2013, ApJ, 763, L32
Chauvin, G., Vigan, A., Bonnefoy, M., et al. 2015, A&A, 573, A127
Cumming, A., Butler, R. P., Marcy, G. W., et al. 2008, PASP, 120, 531
Cushing, M. C., Marley, M. S., Saumon, D., et al. 2008, ApJ, 678, 1372
Fuhrmann, K., & Chini, R. 2015, ApJ, 806, 163
Groff, T. D., Kasdin, N. J., Limbach, M. A., et al. 2014, in Proc. SPIE, Vol. 9147, 1
Hinkley, S., Oppenheimer, B. R., Zimmerman, N., et al. 2011, PASP, 123, 74
Hinkley, S., Pueyo, L., Faherty, J. K., et al. 2013, ApJ, 779, 153
Høg, E., Fabricius, C., Makarov, V. V., et al. 2000, A&A, 355, L27
Ingraham, P., Marley, M. S., Saumon, D., et al. 2014, ApJ, 794, L15
Janson, M., Brandt, T. D., Kuzuhara, M., et al. 2013, ApJ, 778, L4
Knapp, G. R., Leggett, S. K., Fan, X., et al. 2004, AJ, 127, 3553
Konopacky, Q. M., Barman, T. S., Macintosh, B. A., & Marois, C. 2013, Science, 339, 1398
Kuzuhara, M., Tamura, M., Kudo, T., et al. 2013, ApJ, 774, 11
Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, ApJ, 660, 770
Lagrange, A.-M., Bonnefoy, M., Chauvin, G., et al. 2010, Science, 329, 57
Lagrange, A.-M., Gratadour, D., Chauvin, G., et al. 2009, A&A, 493, L21
Lee, J.-M., Heng, K., & Irwin, P. G. J. 2013, ApJ, 778, 97
Line, M. R., Teske, J., Burningham, B., Fortney, J. J., & Marley, M. S. 2015, ApJ, 807, 183
Line, M. R., Wolf, A. S., Zhang, X., et al. 2013, ApJ, 775, 137
Macintosh, B., Graham, J. R., Barman, T., et al. 2015, Science, 350, 64
Macintosh, B. A., Graham, J. R., Palmer, D. W., et al. 2008, in Proc. SPIE, Vol. 7015
Madhusudhan, N. 2012, ApJ, 758, 36
Madhusudhan, N., Burrows, A., & Currie, T. 2011, ApJ, 737, 34
Maire, A.-L., Boccaletti, A., Rameau, J., et al. 2014, A&A, 566, A126
Maire, J., Perrin, M. D., Doyon, R., et al. 2010, in Proc. SPIE, Vol. 7735, 31
Mamajek, E. E., & Bell, C. P. M. 2014, MNRAS, 445, 2169
Marley, M. S., Fortney, J. J., Hubickyj, O., Bodenheimer, P., & Lissauer, J. J. 2007, ApJ, 655, 541
Marois, C., Correia, C., Galicher, R., et al. 2014, in Proc. SPIE, Vol. 9148, 0
Marois, C., Doyon, R., Racine, R., & Nadeau, D. 2000, PASP, 112, 91
Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, ApJ, 641, 556
Marois, C., Macintosh, B., Barman, T., et al. 2008, Science, 322, 1348
Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, Nature, 468, 1080

Nielsen, E. L., Liu, M. C., Wahhaj, Z., et al. 2013, ArXiv e-prints, 1306.1233
Öberg, K. I., Murray-Clay, R., & Bergin, E. A. 2011, ApJ, 743, L16
Perrin, M. D., Maire, J., Ingraham, P., et al. 2014, in Proc. SPIE, Vol. 9147, 3
Peters, M. A., Groff, T., Kasdin, N. J., et al. 2012, in Proc. SPIE, Vol. 8446
Peters-Limbach, M. A., Groff, T. D., Kasdin, N. J., et al. 2013, in Proc. SPIE, Vol. 8864, 1
Pueyo, L. 2016, ArXiv e-prints, 1604.06097
Quanz, S. P., Meyer, M. R., Kenworthy, M. A., et al. 2010, ApJ, 722, L49
Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, PASP, 111, 587
Rameau, J., Chauvin, G., Lagrange, A.-M., et al. 2013, ApJ, 772, L15
Rice, E. L., Barman, T., Mclean, I. S., Prato, L., & Kirkpatrick, J. D. 2010, ApJS, 186, 63
Rice, E. L., Oppenheimer, R., Zimmerman, N., Roberts, L. C., & Hinkley, S. 2015, PASP, 127, 479
Saumon, D., & Marley, M. S. 2008, ApJ, 689, 1327
Saumon, D., Marley, M. S., Cushing, M. C., et al. 2006, ApJ, 647, 552
Schneider, G., & Silverstone, M. D. 2003, in Proc. SPIE, ed. A. B. Schultz, Vol. 4860, 1
Skemer, A. J., Marley, M. S., Hinz, P. M., et al. 2014, ApJ, 792, 17
Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJ, 755, L28
Sparks, W. B., & Ford, H. C. 2002, ApJ, 578, 543
Spergel, D., Gehrels, N., Baltay, C., et al. 2015, ArXiv e-prints, 1503.03757
Spiegel, D. S., & Burrows, A. 2012, ApJ, 745, 174
Thalmann, C., Carson, J., Janson, M., et al. 2009, ApJ, 707, L123
Wagner, K., Apai, D., Kasper, M., et al. 2016, Science, 353, 673
Wahhaj, Z., Liu, M. C., Biller, B. A., et al. 2013, ApJ, 779, 80